Course notes, CS664, 9/7/04

MOTION

- Motion problem definition: like stereo, but temporally separated images.

- Can be viewed as very similar to stereo, just with 2D disparities (shifts).

- Almost all motion algorithms work for stereo and vice-versa, just nowhere near as well. Visual correspondence is just code for "motion and stereo".

- A natural problem to think about in a continuous framework, as if we are sampling from $I(x, y, t)$ at closely spaced intervals in t.

- This leads to some very elegant math that doesn't work particularly well in practice. But it's a good thing to know about, so we will cover it.

- Some nice statistical techniques that can be used both for motion and for stereo were first developed for motion.

- Note: under some circumstances the MR "camera" can directly yield motion (in a single image).

- Some applications of motion work directly on the output of motion. Examples: detecting cyclical motion patters, such as a person walking.

- Another example: detecting the focus of expansion/contraction, as in "Star Trek". Useful for robots (can infer direction of heading).

- Also consider motion segmentation done with a static camera, i.e. subtraction. Ideally you would like to do "event detection".

- There is some very nice work on this done at TI (elegant but simple). Based on James Courtney's 1997 paper. Even though it just uses subtraction, the ideas it layers on top are quite nice

(anyone know how to do this with a moving camera?)

GENERALIZATIONS OF MOTION AND STEREO

- Next topic: three generalizations of motion and stereo
  1. Registration
  2. Image restoration, and pixel-labeling problems
  3. Feature-space analysis

REGISTRATION

- The registration problem is a variant of the correspondence problem.

- Think about trying to two images with some overlap (like snapshots). You might want to align them so that their overlaps "mesh", thus giving you a bigger image.

- Original application: aerial surveillance, for forestry or spying (or both!)

- Two current applications: constructing mosaics (panoramas) and medical imaging

- Medical imaging: suppose you have two images taken with different modalities (such as MR and CT). Different boundaries might be visible from different modalities. A common use of this is to align one entire image so it overlaps correctly with the other.

- Algorithms essentially rely on a similarity measure between images. Consider a set of transformations T; we want to find

$$\arg\min_{t \in T} |I_1 - t(I_2)|.$$

- What are T and $|\cdot|$?

- T is a transformation, mapping from pixels to pixels. Simplest one is just a 2-dof translation, but there are lots of variants that include rotations, scaling, etc. As a rule, the more complex T is the more accurate it is but the slower the search is.

- With the right camera geometry T really is just translations.

Rotations about the optical center. Under these circumstances, with a stationary scene, pixels in the (overlapping portion of the) two images are actually *identical*, up to sampling. This is cool but only practical for certain applications (anyone remember Quicktime VR?) Used for certain web "walkthroughs".

- Image similarity measures have lots of other uses (such as image search). We will see a few algorithms (correlation, color histograms, mutual information, distance transforms)